

20

---

## A New Model for Probabilistic Information Retrieval on the Web

**GABRIEL MATEESCU**

Research Computing Support Group,  
National Research Council Canada,  
Ottawa, Canada

**MASHA SOSONKINA**

Scalable Computing Laboratory,  
Ames Laboratory, Ames, USA

**PAUL THOMPSON**

Thayer School of Engineering,  
Dartmouth College, Hanover, USA

### ABSTRACT\*

Academic research in information retrieval did not make its way into commercial retrieval products until the last 15 years.

---

\*Support for Paul Thompson for this research was provided by a Department of Defense Critical Infrastructure Protection Fellowship grant with the Air Force Office of Scientific Research, F49620-01-1-0272.

Early web search engines also made little use of information retrieval research, in part because of significant differences in the retrieval environment on the Web, such as higher transaction volume and much shorter queries. Recently, however, academic research has taken root in search engines. This paper describes recent developments with a probabilistic retrieval model originating prior to the Web, but with features which could lead to effective retrieval on the Web. Just as graph structure algorithms make use of the graph structure of hyperlinking on the Web, which can be considered a form of relevance judgments, the model of this paper suggests how relevance judgments of web searchers, not just web authors, can be taken into account in ranking. This paper also shows how the combination of expert opinion probabilistic information retrieval model can be made computationally efficient through a new derivation of the mean and standard deviation of the model's main probability distribution.

## 1. INTRODUCTION

Academic research in information retrieval did not make its way into commercial retrieval until the last 15 years when products such as Personal Librarian (Koll, 1981) or, later the ranked retrieval mode of West-law, WIN (West Publishing Company, 1993) became available. Early web search engines also made little use of information retrieval research, in part because of significant differences in the retrieval environment on the Web. Two main differences from earlier retrieval paradigms include higher transaction volume and much shorter queries. More recently, however, academic research has taken root in search engines such as Google (Brin et al., 1998).

This paper describes recent developments with a probabilistic retrieval model that originated prior to the Web, but which has features that may lead to effective retrieval on the Web. Just as graph structure algorithms make use of the graph structure of hyperlinking on the Web (see, for example, Brin et al., 1998, Kleinberg, 1999), which can be considered a form of relevance judgments, the model of this paper

shows how the relevance judgments of web searchers, not just web authors, can be taken into account in ranking. This paper also shows how the combination of expert opinion probabilistic information retrieval model can be made computationally efficient through a new derivation of the mean and standard deviation of the model's main probability distribution.

## 2. BACKGROUND

The Bayesian Combination of Expert Opinion (CEO) approach to probabilistic information retrieval was first described by Thompson (1986, 1990a,b). The CEO model is a generalization of the unified probabilistic retrieval model developed by Robertson et al. (1982), also known as the RMC model. The unified model, called Model 3 in Robertson et al. (1982), was an attempt to combine the models of probabilistic information retrieval developed by Maron and Kuhns (1960), referred to as Model 1, with the probabilistic retrieval model developed by Robertson and Sparck Jones (1976), van Rijsbergen (1979), Croft and Harper (1979), and others, referred to as Model 2. As it has been the case with most probabilistic retrieval models, these models were based on the use of point probabilities, rather than on probability distributions.

The CEO model, by contrast, provides a probability distribution for a document's being judged relevant by a particular user. Both the mean and standard deviation of the distributions are needed in the CEO model for the combination process, as well as the ranking of retrieved documents by probability of relevance. In early accounts of the model (Thompson, 1990a,b), it was not shown how the mean and standard deviation, or variance, of these distributions could be computationally implemented. This paper shows how the mean and standard deviation of the CEO model's distribution can be computed and how the CEO model can be applied to Web document retrieval.

The unified probabilistic retrieval model, Model 3, was developed so that probabilistic evidence of relevance from

the two earlier probabilistic models, Models 1 and 2, could be combined in order to produce a more accurate ranking of documents. As stated in the probability ranking principle (van Rijsbergen, 1979):

If a reference retrieval system's response to each request is a ranking of the documents in the collection in order of decreasing probability of relevance to the user who submitted the request, where the probabilities are estimated as accurately as possible on the basis of whatever data have been made available to the system for this purpose, the overall effectiveness of the system to its user will be the best that is obtainable on the basis of those data.

There were several unresolved issues with the RMC version of Model 3. Robertson (1984) has shown that the term independence assumptions on which the model is based lead to inconsistencies. Moreover, the RMC version of Model 3 did not support relevance feedback. The CEO model, which is based on Model 3, was developed to overcome these difficulties, as well as to provide a general probabilistic retrieval model that could combine probabilities from multiple probabilistic retrieval models, not only the two models unified by the RMC model.

In particular, it explored the use of subjective probabilities provided by indexers or searchers (Thompson, 1988).

### 3. RELATED RESEARCH

The past decade has seen much research on the combination of results from multiple retrieval algorithms, representations of text and query, and retrieval systems (Croft, 2000). The motivation for this research has been provided by several empirical studies showing that different algorithms, representations, and systems provide substantially different, though overlapping, sets of relevant documents (Croft and Harper, 1979; Katzer et al., 1982; McGill et al., 1979; Saracevic and Kantor, 1988). This activity has manifested itself both in academic research and in the commercial development of various Web metasearch engines (Aslam and

1 Montague, 2001; Belkin et al., 1995; Manmatha et al., 2001;  
2 Selberg and Etzioni, 1999; Selberg, 1999).

3 Combination of models has also been an active area of  
4 research in other fields, including statistics (Hoeting et al.,  
5 1999; Moerland, 1999, 2000), statistical decision theory  
6 (Clemen and Winkler, 1999; Lindley, 1983; Roback and  
7 Givens, 2001), statistical pattern recognition (Jain et al.,  
8 2000; Xu et al., 1992), machine learning (Freund and  
9 Schapire, 1997; Littlestone and Warmuth, 1992; Lewis et al.,  
10 1996; Schapire and Singer, 1998) and neural networks  
11 (Hashem et al., 1994, 1997; Hofmann and Puzhica, 1998;  
12 Hofmann et al., 1999; Jordan and Jacobs, 1994; Jacobs et al.,  
13 1991; Tumer and Ghosh, 1996, 1999).

14 Many researchers have applied machine learning tech-  
15 niques to automatic text categorization or clustering, see, for  
16 example, Lewis et al. (1996). Mathematical techniques new  
17 to document retrieval, such as singular value decomposition,  
18 or latent semantic indexing, have also been applied  
19 (Deerwester, 1990). More recently, probabilistic variants of  
20 latent semantic indexing have been implemented as well  
21 (Hofmann, 1999, 2001; Papadimitriou et al., 1998).

#### 22 23 24 **4. THE COMBINATION OF EXPERT OPINION** 25 **MODEL**

26  
27 The CEO model is a version of Model 3 that uses probability  
28 distributions, while the RMC version uses point probabilities.  
29 Furthermore, unlike the RMC model, which is based on point  
30 reconciliation, the CEO model applies Lindley's approach  
31 (Lindley, 1983) to reconciliation of probability distributions  
32 to probabilistic information retrieval. In this Bayesian model,  
33 a decision maker with an initial, or prior, probability (or dis-  
34 tribution) for some event or parameter  $\theta$ , consults  $n$  experts  
35 who provide their probabilities (or distributions) as evidence  
36 with which to update the decision maker's prior probability  
37 distribution via Bayes' theorem to obtain a revised, or poster-  
38 ior, probability (or distribution). In the CEO approach,  
39 there are two levels of combination. At the upper level, the

probabilistic information retrieval system is considered the decision maker, and Models 1 and 2 the experts. At the lower level, Models 1 and 2 themselves are derived from CEO. The indexer, or user, in Model 1, or 2, respectively, is seen as a multiple expert—an expert with respect to each use or document property. Each expert, or decision maker, is estimating  $\theta^*$ , the chance of relevance of a document with respect to a query, i.e., the long-run relative frequency of success in a Bernoulli sequence of relevance judgments of users for documents. Each Bernoulli sequence is different, but there is a common subsequence that underlies each, so that each expert, or decision maker, can be seen as estimating  $\theta^*$ , for the underlying subsequence. The parameter actually used in the model is  $\theta$ , the log-odds function of  $\theta^*$ , i.e.,  $\theta = \log[\theta^*/(1-\theta^*)]$ .

In the CEO model, the evidence provided to the decision maker (or information system) by the models being combined is the set of mean values and standard deviations of the distributions provided by the experts consulted by the decision maker. Let  $p(m|s, \theta)$  denote the decision maker's probability distribution for the expert saying that the mean for the log-odds of the chance of relevance is  $m$ , given that the expert provides the standard deviation  $s$  for the log-odds of the chance of relevance and given the true value of  $\theta$ . The decision maker's opinion of the experts' expertise, i.e., the weighting of the experts' evidence, is expressed by assuming that  $p(m|s, \theta)$  is normal with mean  $\alpha + \beta\theta$  and standard deviation  $\gamma s$ , where  $\alpha$ ,  $\beta$  and  $\gamma$  are parameters that can be determined either through data (relevance judgments) or a decision maker's subjective belief ( $\gamma s$  is the result of modifying  $s$  by the factor  $\gamma$ , as called for in Lindley's model for reconciliation (Lindley et al., 1983), on which the CEO model is based).

A simplified version of the CEO model was used in the first and second Text Retrieval Conferences (TREC) (Thompson, 1993, 1994). In this version, the mean and standard deviation of the model's main distribution were calculated using approximate techniques. More importantly, relevance feedback was not incorporated in TREC 1. In TREC

2, a form of relevance feedback was used. The ranked retrieval models combined in the TREC 1 system were weighted by their performance. Unfortunately, due to the many changes made to the models between TREC 1 and TREC 2, the models' performance on TREC 2 was not well predicted by their Model 1 performance.

## 5. RELEVANCE FEEDBACK

Relevance feedback, i.e., the incorporation of users' judgments as to the relevance of retrieved documents to their information needs, presented a problem with pre-web retrieval. Laboratory experiments showed that large gains in performance, in terms of precision and recall, could be gained through use of relevance feedback (Ide and Salton, 1971). On the other hand, it was assumed that it would not be possible to induce users to provide relevance judgments. Westlaw's WIN was introduced without a relevance feedback capability (West Publishing, Company, 1993). By contrast, Lexis-Nexis' Freestyle and some web search engines introduced commands that provided relevance feedback based on a single document, rather than a set of relevant documents. These are often called "more like this" commands, where a user selects a single highly relevant returned document and the system returns similar documents. In the TREC conferences and other experimental settings, use has been made of pseudo-relevance feedback, where the top  $n$  documents are assumed to be relevant and relevance feedback is calculated as though these  $n$  documents had actually been judged relevant (Sakai et al., 2001). As pointed out by Croft et al. (2001), early work on relevance feedback was done with collections of abstracts and results with full text documents have not been as good as was anticipated.

In addition to this type of more or less traditional relevance feedback, new forms of relevance feedback have emerged, including implicit relevance feedback, e.g., systems such as Direct Hit (DH) (which provides relevance feedback based on mining a user's clickstream), recommender systems (Herlocker, 2001), and rating systems (Dellarocas, 2001).

Relevance feedback is usually seen as taking place during a single user's search, but relevance feedback has also been considered in more persistent ways, e.g., in dynamic document spaces (Brauen, 1971). In dynamic document spaces, a user's relevance judgments permanently modify the weights of index terms associated with documents.

## 6. PROBABILITY DISTRIBUTIONS IN THE COMBINATION OF EXPERT OPINION

As mentioned above, each probabilistic model, e.g., the indexer or the user, is making an estimate of  $\theta$  for a common underlying Bernoulli subsequence of the overall Bernoulli sequence of viewings of documents by users. Because each model is making these judgments based on the conditioning information available to it, that model's judgment for the sequence's distribution is exchangeable, i.e., the distribution is invariant under finite permutations of its indices (de Finetti, 1974). A natural distribution to use for a parameter that ranges from 0 to 1, e.g., the proportion of successes in a sequence of relevance judgments, is the beta distribution (Bunn, 1984). It can be very simply updated with each relevance judgment. Graphically, the beta distribution can take many shapes, and is thus capable of expressing a wide range of opinions. The CEO algorithm uses a transformation of the beta distribution, the distribution of  $\log[x/(1-x)]$ , where  $x$  is a random variable with a beta distribution. It is this distribution, referred to as the transformed beta distribution, from which the mean and standard deviation need to be extracted in order to perform the combination of expert opinion and to probabilistically rank retrieved documents.

## 7. COMPUTING THE MEAN AND STANDARD DEVIATION OF THE TRANSFORMED BETA DISTRIBUTION

Let  $y$  be a continuous random variable whose distribution function is the transformed beta distribution. In this section,



we derive expressions for the mean value and standard deviation of  $y$ .

The moment generating function of  $y$  is the function  $\psi$  defined by (see Thompson, 1990b):

$$\psi(t) = \frac{\Gamma(p+t)\Gamma(q-t)}{\Gamma(p)\Gamma(q)}$$

where  $p, q > 0$ ,  $-\infty < x < \infty$  and  $\Gamma(x)$  is the Gamma function defined by

$$\Gamma(x) = \int_0^{\infty} t^{x-1} e^{-t} dt$$

### 7.1. Computation of the Mean Value

The mean value  $\mu$  of the random variable  $y$  is derived as follows

$$\begin{aligned} \mu &= \left. \frac{d\psi(t)}{dt} \right|_{t=0} = \left. \frac{d}{dt} \frac{\Gamma(p+t)\Gamma(q-t)}{\Gamma(p)\Gamma(q)} \right|_{t=0} \\ &= \left. \frac{d}{dt} \frac{\Gamma(p+t)}{\Gamma(p)} \right|_{t=0} \left. \frac{\Gamma(q-t)}{\Gamma(q)} \right|_{t=0} + \left. \frac{d}{dt} \frac{\Gamma(q-t)}{\Gamma(q)} \right|_{t=0} \left. \frac{\Gamma(p+t)}{\Gamma(p)} \right|_{t=0} \\ &= \left. \frac{\Gamma'(p+t)}{\Gamma(p)} \right|_{t=0} - \left. \frac{\Gamma'(q-t)}{\Gamma(q)} \right|_{t=0} \end{aligned} \quad (1)$$

where we have used these facts

$$\left. \frac{d\Gamma(p+t)}{dt} \right|_{t=0} = \Gamma'(p) \quad \text{and} \quad \left. \frac{d\Gamma(q-t)}{dt} \right|_{t=0} = -\Gamma'(q) \quad (2)$$

Because  $\Gamma'(x \pm t)|_{t=0} = \Gamma'(x)$  we get from (1)

$$\mu = \frac{\Gamma'(p)}{\Gamma(p)} - \frac{\Gamma'(q)}{\Gamma(q)} \quad (3)$$

It can be shown (see Whittaker and Watson, 1990) that

$$\Gamma(x) \lim_{n \rightarrow \infty} \frac{n^n n!}{x(x+1) \cdots (x+n)} \quad (4)$$

and it is easy to show that

$$\frac{n^x n!}{x(x+1) \cdots (x+n)} = e^{x(\ln n - 1 - (1/2) - \cdots - (1/n))} \frac{1}{x} \frac{e^{x/1}}{1 + (x/1)} \frac{e^{x/2}}{1 + (x/2)} \cdots \frac{e^{x/n}}{1 + (x/n)}$$

Therefore, substituting in (4) we get

$$\Gamma(x) = e^{-Cx} \frac{1}{x} \prod_{n=1}^{\infty} \frac{e^{x/n}}{1 + (x/n)} \quad (5)$$

where  $C$  is the Euler–Macheroni constant defined by the limit

$$C = \lim_{x \rightarrow \infty} \left( 1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{n} - \ln n \right)$$

and the value of  $C$  computed with 10 decimal places is

$$C = 0.5772156649$$

Taking the logarithm of (5) and differentiating gives:

$$\frac{\Gamma'(x)}{\Gamma(x)} = -C - \frac{1}{x} + \sum_{i=1}^{\infty} \frac{x}{i(x+i)} \quad (6)$$

We are now ready to compute  $\mu$  from (3):

$$\begin{aligned} \mu &= -C - \frac{1}{p} + \sum_{i=1}^{\infty} \frac{p}{i(p+i)} - \left( -C - \frac{1}{q} + \sum_{i=1}^{\infty} \frac{q}{i(q+i)} \right) \\ &= \frac{p-q}{pq} + \sum_{i=1}^{\infty} \frac{p}{i(p+i)} - \sum_{i=1}^{\infty} \frac{q}{i(q+i)} \end{aligned} \quad (7)$$

## 7.2. Computation of the Standard Deviation

The standard deviation  $\sigma^2$  of  $y$  is defined as

$$\sigma^2 = E[y^2] - \mu^2$$

with  $E[y^2]$  being the second moment of  $y$  that is computed using the formula:

$$\begin{aligned}
 E[y^2] &= \left. \frac{d^2 \psi(t)}{dt^2} \right|_{t=0} = \left. \frac{d^2 \Gamma(p+t)\Gamma(q-t)}{dt^2 \Gamma(p)\Gamma(q)} \right|_{t=0} \\
 &= \left. \frac{\Gamma''(p+t)\Gamma(q-t) - 2\Gamma'(p+t)\Gamma'(q-t) + \Gamma(p+t)\Gamma''(q-t)}{\Gamma(p)\Gamma(q)} \right|_{t=0}
 \end{aligned}$$

where we have used the facts (2). Thus

$$E[y^2] = \frac{\Gamma''(p)}{\Gamma(p)} + \frac{\Gamma''(q)}{\Gamma(q)} - 2 \frac{\Gamma'(p)\Gamma'(q)}{\Gamma(p)\Gamma(q)}$$

Formally differentiating (6), we get

$$\frac{\Gamma''(x)}{\Gamma(x)} = \left( \frac{\Gamma'(x)}{\Gamma(x)} \right)^2 + \frac{1}{x^2} + \sum_{i=1}^{\infty} \frac{1}{(x+i)^2}$$

then, substituting (6) in the above relation, we obtain:

$$\frac{\Gamma''(x)}{\Gamma(x)} = \left( -C - \frac{1}{x} + \sum_{i=1}^{\infty} \frac{x}{i(x+i)} \right)^2 + \frac{1}{x^2} + \sum_{i=1}^{\infty} \frac{1}{(x+i)^2} \quad (8)$$

The second moment of  $y$  is

$$\begin{aligned}
 E[y^2] &= \left( -C - \frac{1}{p} + \sum_{i=1}^{\infty} \frac{p}{i(p+i)} \right)^2 \frac{1}{p^2} + \sum_{i=1}^{\infty} \frac{1}{(p+i)^2} \\
 &\quad + \left( -C - \frac{1}{q} + \sum_{i=1}^{\infty} \frac{q}{i(q+i)} \right)^2 \frac{1}{q^2} + \sum_{i=1}^{\infty} \frac{1}{(q+i)^2} \\
 &\quad - 2 \left( -C - \frac{1}{p} + \sum_{i=1}^{\infty} \frac{p}{i(p+i)} \right) \left( -C - \frac{1}{q} + \sum_{i=1}^{\infty} \frac{q}{i(q+i)} \right) \quad (9)
 \end{aligned}$$

Therefore, the standard deviation and the mean value can be computed approximately by replacing  $\sum_{i=1}^{\infty}$  with  $\sum_{i=1}^n$  in relations (7) and (9), then taking  $n$  large enough to meet a given convergence criterion.

## 8. DISCUSSION

The CEO model provides a probabilistic framework for combining probabilistic retrieval models. The model can be used with subjective probabilities provided, either explicitly or implicitly, by users. It can be used both within the context of a single search and over time. Search on the Web is different in various ways from traditional online document retrieval. Two of the main differences, higher transaction volume and shorter queries, are differences that can be taken advantage of by the CEO model. First, high transaction volumes mean that there are more documents being seen by users from which relevance judgments can be collected. Second, because queries are so much shorter, on average less than three words per query, as compared to seven words more typical of traditional online retrieval, it is important to extend the focus of probabilistic models beyond words in documents and queries. As mentioned above, algorithms such as HITS (Kleinberg, 1999) or Page Rank (Brin et al., 1998) extend the focus to hyperlinking. The CEO model shows how this focus could be further extended to user's relevance judgments, whether explicit or implicit. The statistical model of the reconciliation of probability distributions, on which the CEO algorithm is based, has seen significant development in recent years, e.g., (Roback and Givens, 2001). Related work has been done in machine learning, e.g., on the weighted majority voting algorithm and on boosting (Freund and Schapire, 1997; Littlestone and Warmuth, 1992; Lewis et al., 1996; Schapire and Singer, 1998), mixture models (Cohn and Hofmann, 2001; Hofmann and Puzhica, 1998; Hofmann, 1999, 2001; Hofmann et al., 1999; Jacobs et al., 1991; Jordan and Jacobs, 1994; Moerland, 1999, 2000; Manmatha et al., 2001), and Bayesian model averaging (Hoeting et al., 1999). Text categorization and clustering have become significant application domains for machine learning research. Algorithms such as boosting (Schapire and Singer, 1998) and support vector machines (Joachim, 2001) have achieved good results with text categorization.

1       The focus of these new machine learning and related  
2 techniques has been on the document collection, not on  
3 the user and the user's information need. As noted by  
4 Papadimitriou et al. (1998), "The approach in this body of work  
5 (probabilistic information retrieval) is to formulate infor-  
6 mation retrieval as a problem of learning the concept of  
7 "relevance" that relates documents and queries. The corpus  
8 and its correlations play no central role. In contrast, our focus  
9 is on the probabilistic properties of the corpus." This focus on  
10 the collection ignores the probabilistic evidence provided by  
11 an analysis of the user and the user's information need.  
12 Relevance is better understood as a relation between the  
13 user's information need, which is represented by the query,  
14 and the intellectual content of the document, which is repre-  
15 sented by the text of the document (Wilson, 1973) . While the  
16 text of queries and documents may model this latent, deeper  
17 structure, especially in the case of the document, user's  
18 relevance judgments (Croft et al., 2001) and mixed-initiative  
19 interaction (Haller et al., 1999) provide additional evidence of  
20 the user's information need. Much research in probabilistic  
21 information retrieval is currently focused on language models  
22 (Callan et al., 2001; Ponte and Croft, 1998; Ponte 1998).  
23 Language models are also mainly applied to collections,  
24 rather than users, though Lafferty and Zhai (2001) provide  
25 two language models, one for the document and one for the  
26 query, and perform retrieval by measuring the similarity of  
27 the two language models.

28       The CEO model predated much of the research dis-  
29 cussed above in the fields of statistics, neural networks,  
30 and machine learning. Lindley's (1983) model of reconcilia-  
31 tion of distributions, now called Supra-Bayesian pooling, on  
32 which the CEO model is based, is still one of the leading  
33 theories in the Bayesian approach to combining expert  
34 opinions (Roback and Givens, 2001). The basic framework  
35 of the CEO model appears to be sound, but the model still  
36 needs to be completely implemented and empirically tested.  
37 In the process of doing so it is likely that the model can be  
38 improved through the incorporation of some aspects of the  
39 more recent research discussed above. In particular,

although there is long-standing precedence in the decision theory literature (Bunn, 1984) for using the beta distribution, as discussed above, to model expert opinion, it may be that techniques from Bayesian model averaging (Hoeting et al., 1999) could lead to more accurate modeling. With respect to representation of experts' opinion, the CEO model only requires a mean and standard deviation, not a specific distributional form. More generally, mixture models now being explored in the context of information retrieval, e.g., (Cohn and Hofmann, 2001; Hofmann, 1999, 2001; Manmatha et al., 2001), may inform new developments with the CEO model.

## 9. CONCLUSION

The probability ranking principle calls for taking all available evidence into account when probabilistically ranking documents in response to a user's request.

The CEO algorithm provides a formalism for taking all such evidence into account using Bayesian subjective decision theory. The theoretical strength of the CEO algorithm, its ability to easily incorporate relevance judgments and use the judgments to continuously tune its probability estimates, has also been its practical weakness. The success of recommender and similar systems in some domains, e.g., e-commerce, shows that implicit relevance judgments can be effective and may lead to settings where algorithms such as CEO, which rely heavily on relevance judgments, can be effective. Now that an efficient method of calculating the mean and standard deviation of the transformed data distribution has been derived, the implementation of the CEO model will be facilitated.

## REFERENCES

- Aslam, J. A., Montague, M. (2001). Models for metasearch. In: Croft, W. B., Harper, D. J., Kraft, D., Zobel, J., eds. *Proceedings of the 24th Annual Retrieval (SIGIR' 01)*. New Orleans, Sept. 9–13, ACM Press, pp. 276–284.

- 1 Belkin, N., Kantor, P., Fox, E., Shaw, J. (1995). Combining the  
2 evidence of multiple query representations for information  
3 retrieval. *Information Processing & Management* 31(3):  
4 431–448.
- 5 Brauen, T. L. (1971). Document vector modification. Salton, G., ed.  
6 *Experiments in Automatic Document Processing*. Englewood  
7 Cliffs, NJ: Prentice-Hall, 456–484.
- 8 Brin, S., Motwani, R., Page, L., Winograd, T. (1998). What can you  
9 do with a web in your pocket? *Bulletin of the IEEE Computer*  
10 *Society Technical Committee on Data Engineering*. AQ1
- 11 Bunn, D. W (1984). *Applied Decision Analysis*. New York, NY:  
12 McGraw-Hill.
- 13 Callan, J., Croft, B., Lafferty, J. (2001). Proceedings of the  
14 Workshop on Language Modeling and Information Retrieval.  
15 Carnegie Mellon University, May 31–June 1. <http://la.lti.cs.cmu.edu/callan/Workshops/lmir01/WorkshopProcs>.
- 16 Clemen, R. T., Winkler, R. L. (1999). Combining probability dis-  
17 tributions from experts in risk analysis. *Risk Analysis* 19:  
18 187–203.
- 19 Cohn, D., Hofmann, T. (2001). The missing link: a probabilistic  
20 model of document content and hypertext connectivity.  
21 *Advances in Neural Information Processing Systems*  
22 *(NIPS\* 13)*. Cambridge, MA: MIT Press.
- 23 Croft, W. B. (2000). Combining approaches to information retrieval.  
24 Croft, W. B., ed. *Advances in Information Retrieval: Recent*  
25 *Research from the Center for Intelligent Information Retrieval*.  
26 Chapter 1. Boston, MA: Kluwer Academic, 1–36.
- 27 Croft, W. B., Harper, D. J. (1979). Using probabilistic models of  
28 document retrieval without relevance information. *Journal*  
29 *of Documentation* 45(4):285–295.
- 30 Croft, W. B., Cronen-Townsend, S., Lavrenko, V. (2001). Relevance  
31 feedback and personalization: a language modeling pers-  
32 pective. In: Joint DELOS-NSF Workshop on Personalization  
33 and Recommender Systems in Digital Libraries. Dublin,  
34 Ireland, June 18–20, pp. 18–20.
- 35 Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K.,  
36 Harshman, R. (1990). Indexing by latent semantic analysis.

- 1        *Journal of the American Society for Information Science*  
 2        41(6):381–407.
- 3        Dellarocas, C. (2001). Building trust on-line: the design of reliable  
 4        reputation reporting mechanisms for online trading commu-  
 5        nities. Paper 101. Center for eBusiness@MIT.
- 6        de Finetti, B. (1974) (1974). *Theory of Probability: A Critical*  
 7        *Introductory Treatment*. Vol. 1. New York: John Wiley.
- 8        Direct Hit. <http://www.directhit.com>. AQ2
- 9        Freund, Y., Schapire, R. E. (1997). A decision-theoretic generaliza-  
 10        tion of on-line learning and an application to boosting. *Journal*  
 11        *of Computer and System Science*. 55(1): 119–139.
- 12        Haller, S., McRoy, S., Kobsa, A. (1999) (1999). *Computational Models*  
 13        *of Mixed-Initiative Interaction*. Boston, MA: Kluwer Academic.
- 14        Hashem, S. (1997). Algorithms for optimal linear combinations  
 15        of neural networks. *International Conference on Neural*  
 16        *Networks* Vol. 1, pp. 242–247.
- 17        Hashem, S., Schmeiser, B., Yih, Y. (1994). Optimal linear com-  
 18        binations of neural networks: an overview. IEEE World  
 19        Congress on Computational Intelligence, IEEE International  
 20        Conference on Neural Networks, Vol. 3, pp. 1507–1512.
- 21        Herlocker, J. (2001). ACM SIGIR Workshop on Recommender  
 22        Systems Notes. [http://www.cs.orst.edu/~herlock/rsw2001/](http://www.cs.orst.edu/~herlock/rsw2001/workshop_notes.html)  
 23        [workshop\\_notes.html](http://www.cs.orst.edu/~herlock/rsw2001/workshop_notes.html).
- 24        Hoeting, J. A., Madigan, D., Raftery, A. E., Volinsky, C. T. (1999).  
 25        Bayesian model averaging: a tutorial. *Statistical Science*  
 26        14(4):382–417.
- 27        Hofmann, T. (1999). Probabilistic latent semantic indexing. In:  
 28        Hearst, M., Gey, F., Tong, R., eds. Proceedings of the 22nd  
 29        Annual International ACM SIGIR Conference on Research  
 30        and Development in Information Retrieval (SIGIR' 99),  
 31        Berkeley, CA, ACM Press, Aug. 15–19, pp. 50–57.
- 32        Hofmann, T. (2001). What people (don't) want. European Con-  
 33        ference on Machine Learning (ECML).
- 34        Hofmann, T., Puzhica, J. (1998). Unsupervised learning from dyadic  
 35        data. Technical Report TR-98-042, University of California,  
 36        Berkeley.



- 1 Hofmann, T., Puzhica, J., Jordan, M. (1999). Learning from  
2 dyadic data. In: Kearns, M. S., Solla, S. A., Cohn, D., eds.  
3 *Advances in Neural Information Processing Systems*.  
4 Number 11. Cambridge, MA: MIT Press.
- 5 Ide, E., Salton, G. (1971). Interactive search strategies and dynamic  
6 file organization in information retrieval. In: Salton, G., ed.  
7 *Experiments in Automatic Document Processing*. Englewood  
8 Cliffs, NJ: Prentice-Hall, pp. 373–393.
- 9 Jacobs, R. A., Jordan, M. I., Nowlan, S., Hinton, G. E. Adaptive  
10 mixtures of local experts. *Neural Computation* 3:1–12.
- 11 Jain, A. K., Duin, R. P. W., Mao, J. (2000). Statistical pattern recog-  
12 nition: a review. *IEEE Transactions on Pattern Analysis and*  
13 *Machine Intelligence* 22(1):4–37.
- 14 Joachim, T. (2001). A statistical learning model of text categoriza-  
15 tion for support vector machines. In: Croft, W. B., Harper, D. J.,  
16 Kraft, D., Zobel, J., eds. Proceedings of the 24th Annual Inter-  
17 national ACM SIGIR Conference on Research and Develop-  
18 ment in Information Retrieval (SIGIR' 01), New Orleans,  
19 Sept. 9–13, ACM Press, pp. 128–136.
- 20 Jordan, M. I., Jacobs, R. A. (1994). Hierarchical mixtures of experts  
21 and the EM algorithm. *Neural Computation* 6:181–214.
- 22 Katzer, J., McGill, M., Tessier, J., Frakes, W., Das-Gupta, P. (1982).  
23 A study of the overlap among document representations infor-  
24 mation technology. *Research and Development* 1(4):261–274.
- 25 Kleinberg, J. M. (1999). Authoritative sources in a hyperlinked  
26 environment. *Journal of the ACM* 46(5):604–622.
- 27 Koll, M. B. (1981). Information retrieval theory and design based on  
28 a model of the user's concept relations. Oddy, R. N., Robertson,  
29 S. E., van Rijsbergen, C. J., Williams, P. W., eds. *Information*  
30 *Retrieval Research*. London: Butterworths.
- 31 Lafferty, J., Zhai, C. (2001). Document language models, query  
32 models, and risk minimization for document retrieval.  
33 In: Croft, W. B., Harper, D. J., Kraft, D., Zobel, J., eds.  
34 Proceedings of the 24th Annual International ACM SIGIR  
35 Conference on Research and Development in Information  
36 Retrieval (SIGIR' 01). New Orleans, Sept. 9–13, ACM Press,  
37 pp. 111–119.
- 38  
39

- 1 Lewis, D. D., Schapire, R. E., Callan, J. P., Papka, R. (1996).  
2 Training algorithms for linear text classifiers. In: Frei, H. P.,  
3 Harman, D., Schauble, P., Wilkinson, R., eds. Proceedings of  
4 the 19th Annual International, ACM SIGIR Conference  
5 on Research and Development in Information Retrieval  
6 (SIGIR' 96). Zurich, Switzerland, Aug. 18–22, ACM Press,  
7 pp. 298–306.
- 8 Lindley, D. V. (1983). Reconciliation of probability distributions.  
9 *Operations Research* 31(5):866–880.
- 10 Littlestone, N., Warmuth, M. (1992). The weighted majority voting  
11 algorithm. Technical Report UCSC-CRL-91-28, University of  
12 California, Santa Cruz.
- 13 Manmatha, R., Rath, T., Feng, F. (2001). Modeling score distri-  
14 butions for combining the outputs of search engines. In: Croft,  
15 W. B., Harper, D. J., Kraft, D., Zobel, J., eds. Proceedings of  
16 the 24th Annual International ACM SIGIR Conference on  
17 Research and Development in Information Retrieval (SIGIR'  
18 01). New Orleans, Sept. 9–13, ACM Press, pp. 267–275.
- 19 Maron, M. E., Kuhns, J. L. (1960). On relevance, probabilistic  
20 indexing and information retrieval. *Journal of the ACM*  
21 7(3):216–244.
- 22 McGill, M., Koll, M., Noreault, T. (1979). An evaluation of factors  
23 affecting document ranking by information retrieval systems.  
24 Final report for grant NSF-ISF-78-10454 to the National  
25 Science Foundation, Syracuse University School of Information  
26 Studies.
- 27 Moerland, P. (1999). A comparison of mixture models for density  
28 estimation. Proceedings of the International Conference on  
29 Artificial Neural Networks (ICANN' 99).
- 30 Moerland, P. (2000). Mixtures of latent variable models for density  
31 estimation and classification. Research Report IDIAP-RR  
32 00-25, IDIAP, Martigny, Switzerland.
- 33 Papadimitriou, C. H., Raghavan, P., Tamaki, H., Vempala, S.  
34 (1998). Latent semantic indexing: a probabilistic analysis.  
35 In: Mendelson, A., Paredaens, J., eds. Proceedings of the  
36 seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on  
37 Principles of Database Systems (PODS' 98). Vol. 3, June 1–4,  
38 Seattle WA, ACM Press, pp. 159–168.
- 39

- Ponte, J. (1998). A language modeling approach to information retrieval. PhD thesis, University of Massachusetts, Amhersts.
- Ponte, J., Croft, W. B. (1998). A language modeling approach to information retrieval. In: Croft, W. B., Moffat, A., van Rijsbergen, C. J., Wilkinson, R., Zobel, J., eds. Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR' 98). Melbourne, Australia, Aug. 24–28, ACM Press, pp. 275–281.
- Roback, P. J., Givens, G. H. (2001). Supra-bayesian pooling of priors linked by a deterministic simulation model. *Communications in Statistics-Simulation and Computation* 30:447–476.
- Robertson, S. E. (1984). Consistency of the RMC model. Personal communication.
- Robertson, S. E., Sparck Jones, K. (1976). Relevance weighting of search terms. *Journal of the American Society for Information Science* 27(3):129–146.
- Robertson, S. E., Maron, M. E., Cooper, W. S. (1982). Probability of relevance: a unification of two competing models for document retrieval. *Information Technology: Research and Development* 1(1):1–21.
- Sakai, T., Robertson, S. E., Walker, S. (2001). Flexible pseudo-relevance feedback for NTCIR-2. In: Proceedings of the Second NTCIR Workshop on Evaluation of Chinese and Japanese Text Retrieval and Text Summarization. Tokyo, Japan. Tokyo: National Institute of Informatics.
- Saracevic, T., Kantor, P. (1988). A study of information seeking and retrieving. III. Searchers, searches, overlap. *Journal of the American Society for Information Science* 39(3):197–216.
- Selberg, E. W. (1999). Towards comprehensive web search. PhD thesis, University of Washington.
- Selberg, E., Etzioni, O. (1997). The MetaCrawler architecture for resource aggregation on the Web. *IEEE Expert* 12(1):8–14.
- Schapire, R. E., Singer, Y. (1998). Improved boosting algorithms using confidence-rated predictions. Proceedings of the Eleventh Annual Conference on Computational Learning Theory, Madison, Wisconsin, July 24–26.

- 1 Thompson, P. (1986). Subjective probability, combination of expert  
2 opinion, and probabilistic approaches to information retrieval.  
3 PhD thesis, University of California, Berkeley.
- 4 Thompson, P. (1988). Subjective probability and information  
5 retrieval: a review of the psychological literature. *Journal of*  
6 *Documentation* 44(2):119–143.
- 7 Thompson, P. (1990a). A combination of expert opinion approach to  
8 probabilistic information retrieval, part 1: The conceptual  
9 model. *Information Processing & Management* 26(3):371–382.
- 10 Thompson, P. (1990b). A combination of expert opinion approach  
11 to probabilistic information retrieval, part 2: Mathematical  
12 treatment of CEO Model 3. *Information Processing & Manage-*  
13 *ment* 26(3):383–394.
- 14 Thompson, P. (1993). Description of the PRC CEO algorithm  
15 for TREC. In: Harman, D. K., ed. The First Text Retrieval  
16 Conference (TREC-1). volume NIST Special Publication  
17 500-207. Washington, DC: US Government Printing Office,  
18 pp. 337–342.
- 19 Thompson, P. (1994) Description of the PRC CEO algorithm for  
20 TREC-2. In: Harman, D. K., ed. The Second Text Retrieval  
21 Conference (TREC-2). volume NIST Special Publication  
22 500-215. Washington, DC: US Government Printing Office,  
23 pp. 271–274.
- 24 Tumer, K., Ghosh, J. (1996). Theoretical foundations of linear  
25 and order statistics combiners for neural pattern classifiers.  
26 Technical Report TR-95-02-98. The University of Texas at  
27 Austin.
- 28 Tumer, K., Ghosh, J. (1999). Linear and order statistics com-  
29 biners for pattern classification. Sharkey, A., ed. *Combining*  
30 *Artificial Neural Networks*. New York, NY: Springer-Verlag,  
31 pp. 127–162.
- 32 van Rijsbergen, C. J (1979). *Information Retrieval*. . 2nd. London,  
33 UK: Buttersworth.
- 34 West Publishing Company. (1993). *WESTLAW Reference Manual*.  
35 5th ed. St. Paul, MN: West Publishing Co.
- 36 Whittaker, E. T., Watson, G. N. (). *A Course in Modern Analysis*.  
37 4th. Cambridge, UK: Cambridge University Press.

- 1 Wilson, P. (1973). Situational relevance. *Information Storage and*  
2 *Retrieval* 9(7):457–471.
- 3 Xu, L., Krzyzak, A., Suen, C. Y. (1992). Methods of combining  
4 classifiers and their applications to handwriting recognition.  
5 *IEEE Transactions on Systems, Man, and Cybernetics*  
6 22(3):418–435.